arXiv:1411.4732v4 [cs.IT] 4 Mar 2015

Quantifying Redundant Information in Predicting a Target Random Variable

Virgil Griffith[1] and Tracey Ho[2]

[1]Computation and Neural Systems, Caltech, Pasadena, CA 91125
[2]Computer Science and Electrical Engineering, Caltech, Pasadena, CA 91125

***Abstract*—This paper considers the problem of defining a measure of redundant information that quantifies how much common information two or more random variables specify about a target random variable. We discussed desired properties of such a measure, and propose new measures with some desirable properties.**

I. Introduction

Many molecular and neurological systems involve multiple interacting factors affecting an outcome synergistically and/or redundantly. Attempts to shed light on issues such as population coding in neurons, or genetic contribution to a phenotype (e.g. eye-color), have motivated various proposals to leverage principled information-theoretic measures for quantifying informational synergy and redundancy, e.g. [1]–[5]. In these settings, we are concerned with the statistics of how two (or more) random variables X_1, X_2, called predictors, jointly or separately specify/predict another random variable Y, called a target random variable. This focus on a target random variable is in contrast to Shannon's mutual information which quantifies statistical dependence between two random variables, and various notions of common information, e.g. [6]–[8].

The concepts of synergy and redundancy are based on several intuitive notions, e.g., positive informational synergy indicates that X_1 and X_2 act cooperatively or antagonistically to influence Y; positive redundancy indicates there is an aspect of Y that X_1 and X_2 can each separately predict. However, it has proven challenging [9]–[12] to come up with precise information-theoretic definitions of synergy and redundancy that are consistent with all intuitively desired properties.

II. Background: Partial Information Decomposition

The *Partial Information Decomposition* (PID) approach of [13] defines the concepts of synergistic, redundant and unique information in terms of *intersection information*, $\mathrm{I}_\cap\big(\{X_1,\ldots,X_n\}:Y\big)$, which quantifies the common information that each of the n predictors $X_1,\ldots,X_n$ conveys about a target random variable Y. An antichain lattice of redundant, unique, and synergistic partial informations is built from the intersection information.

Partial information diagrams (PI-diagrams) extend Venn diagrams to represent synergy. A PI-diagram is composed of nonnegative *partial information regions* (PI-regions). Unlike the standard Venn entropy diagram in which the sum of all regions is the joint entropy $\mathrm{H}(X_{1\ldots n}, Y)$, in PI-diagrams the sum of all regions (i.e. the space of the PI-diagram) is the mutual information $\mathrm{I}(X_{1\ldots n}:Y)$. PI-diagrams show how the mutual information $\mathrm{I}(X_{1\ldots n}:Y)$ is distributed across subsets of the predictors. For example, in the PI-diagram for $n = 2$ (Figure 1): $\{1\}$ denotes the unique information about Y that only X_1 carries (likewise $\{2\}$ denotes the information only X_2 carries); $\{1,2\}$ denotes the redundant information about Y that X_1 as well as X_2 carries, while $\{12\}$ denotes the information about Y that is specified only by X_1 and X_2 synergistically or jointly.

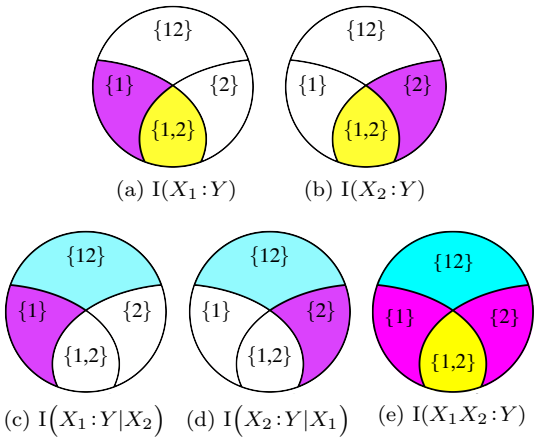

(a) $\mathrm{I}(X_1:Y)$ (b) $\mathrm{I}(X_2:Y)$ (c) $\mathrm{I}\big(X_1:Y|X_2\big)$ (d) $\mathrm{I}\big(X_2:Y|X_1\big)$ (e) $\mathrm{I}(X_1X_2:Y)$

Fig. 1: PI-diagrams for $n = 2$ predictors, showing the amount of redundant (yellow/bottom), unique (magenta/left and right) and synergistic (cyan/top) information with respect to the target Y.

Each PI-region is either redundant, unique, or synergistic, but any combination of positive PI-regions may be possible. Per [13], for two predictors, the four partial informations are defined as follows: the redundant information as

$\mathrm{I}_\cap\big(\{X_1, X_2\}{:}Y\big)$, the unique informations as

$$\begin{aligned}\mathrm{I}_\partial(X_1 : Y) &= \mathrm{I}(X_1{:}Y) - \mathrm{I}_\cap\big(\{X_1, X_2\}{:}Y\big)\\ \mathrm{I}_\partial(X_2 : Y) &= \mathrm{I}(X_2{:}Y) - \mathrm{I}_\cap\big(\{X_1, X_2\}{:}Y\big)\,,\end{aligned} \tag{1}$$

and the synergistic information as

$$\begin{aligned}&\mathrm{I}_\partial(X_1, X_2 : Y)\\ &= \mathrm{I}(X_1, X_2{:}Y) - \mathrm{I}_\partial\left(X_1 : Y\right) - \mathrm{I}_\partial\left(X_2 : Y\right)\\ &\quad - \mathrm{I}_\partial\left(X_1, X_2 : Y\right)\\ &= \mathrm{I}(X_1, X_2{:}Y) - \mathrm{I}(X_1{:}Y) - \mathrm{I}(X_2{:}Y) + \mathrm{I}_\cap\big(\{X_1, X_2\}{:}Y\big)\,.\end{aligned} \tag{2}$$

III. Desired $\mathrm{I}_\cap$ properties and canonical examples

There are a number of intuitive properties, proposed in [5], [9]–[13], that are considered desirable for the intersection information measure $\mathrm{I}_\cap$ to satisfy:

(**$\mathbf{S_0}$**) Weak Symmetry: $\mathrm{I}_\cap\big(\{X_1, \ldots, X_n\}{:}Y\big)$ is invariant under reordering of $X_1, \ldots, X_n$.

(**$\mathbf{M_0}$**) Weak Monotonicity: $\mathrm{I}_\cap\big(\{X_1, \ldots, X_n, Z\}{:}Y\big) \leq \mathrm{I}_\cap\big(\{X_1, \ldots, X_n\}{:}Y\big)$ with equality if there exists $X_i \in \{X_1, \ldots, X_n\}$ such that $\mathrm{H}(Z, X_i) = \mathrm{H}(Z)$.

(**SR**) Self-Redundancy: $\mathrm{I}_\cap\big(\{X_1\}{:}Y\big) = \mathrm{I}(X_1{:}Y)$. The intersection information a single predictor X_1 conveys about the target Y is equal to the mutual information between the X_1 and the target Y.

(**$\mathbf{M_1}$**) Strong Monotonicity: $\mathrm{I}_\cap\big(\{X_1, \ldots, X_n, Z\}{:}Y\big) \leq \mathrm{I}_\cap\big(\{X_1, \ldots, X_n\}{:}Y\big)$ with equality if there exists $X_i \in \{X_1, \ldots, X_n\}$ such that $\mathrm{I}(Z, X_i{:}Y) = \mathrm{I}(Z{:}Y)$.

(**LP**) Local Positivity: For all n, the derived "partial informations" defined in [13] are nonnegative. This is equivalent to requiring that $\mathrm{I}_\cap$ satisfy *total monotonicity*, a stronger form of supermodularity. For $n = 2$ this can be concretized as, $\mathrm{I}_\cap\big(\{X_1, X_2\}{:}Y\big) \geq \mathrm{I}(X_1{:}X_2) - \mathrm{I}(X_1{:}X_2|Y)$.

(**TM**) Target Monotonicity: If $\mathrm{H}\big(Y|Z\big) = 0$, then $\mathrm{I}_\cap\big(\{X_1, \ldots, X_n\}{:}Y\big) \leq \mathrm{I}_\cap\big(\{X_1, \ldots, X_n\}{:}Z\big)$.

There are also a number of canonical examples for which one or more of the partial informations have intuitive values, which are considered desirable for the intersection information measure $\mathrm{I}_\cap$ to attain.

Example Unq, shown in Figure 2, is a canonical case of unique information, in which each predictor carries independent information about the target. Y has four equiprobable states: `ab`, `aB`, `Ab`, and `AB`. X_1 uniquely specifies bit `a`/`A`, and X_2 uniquely specifies bit `b`/`B`. Note that the states are named so as to highlight the two bits of unique information; it is equivalent to choose any four unique names for the four states.

Example RdnXor, shown in Figure 3, is a canonical example of redundancy and synergy coexisting. The `r`/`R` bit is redundant, while the 0/1 bit of Y is synergistically specified as the XOR of the corresponding bits in X_1 and X_2.

X_1	X_2	Y		
a	b	ab	1/4	$\mathrm{I}(X_1, X_2{:}Y) = 2$
a	B	aB	1/4	$\mathrm{I}(X_1{:}Y) = 1$
A	b	Ab	1/4	$\mathrm{I}(X_2{:}Y) = 1$
A	B	AB	1/4	

(a) $\Pr(x_1, x_2, y)$

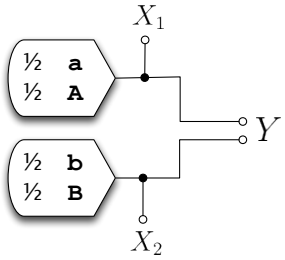

(b) circuit diagram

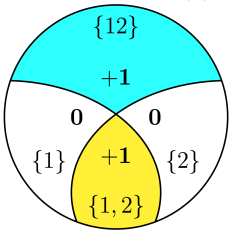

(c) $\mathrm{I}_{\min}$

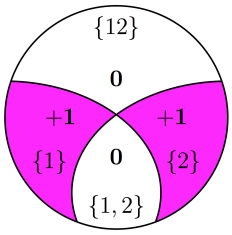

(d) $\mathrm{Syn}/\Delta\mathrm{I}/\mathrm{I}_\wedge/\mathrm{I}_\alpha$

Fig. 2: Example Unq. X_1 and X_2 each uniquely carry one bit of information about Y. $\mathrm{I}(X_1X_2{:}Y) = \mathrm{H}(Y) = 2$ bits.

Example And, shown in Figure 4, is an example where the relationship between X_1, X_2 and Y is nonlinear, making the desired partial information values less intuitively obvious. Nevertheless, it is desired that the partial information values should be nonnegative.

Example ImperfectRdn, shown in Figure 5, is an example of "imperfect" or "lossy" correlation between the predictors, where it is intuitively desirable that the derived redundancy should be positive. Given (**LP**), we can determine the desired decomposition analytically. First, $\mathrm{I}(X_1, X_2{:}Y) = \mathrm{I}(X_1{:}Y) = 1$ bit; therefore, $\mathrm{I}\big(X_2{:}Y|X_1\big) = \mathrm{I}(X_1, X_2{:}Y) - \mathrm{I}(X_1{:}Y) = 0$ bits. This determines two of the partial informations—the synergistic information $\mathrm{I}_\partial(X_1, X_2{:}Y)$ and the unique information $\mathrm{I}_\partial(X_2{:}Y)$ are both zero. Then, the redundant information $\mathrm{I}_\partial(X_1, X_2{:}Y) = \mathrm{I}(X_2{:}Y) - \mathrm{I}_\partial(X_2 : Y) = \mathrm{I}(X_2{:}Y) = 0.99$ bits. Having determined three of the partial informations, we compute the final unique information $\mathrm{I}_\partial(X_1{:}Y) = \mathrm{I}(X_1{:}Y) - 0.99 = 0.01$ bits.

X_1	X_2	Y	
r0	r0	r0	1/8
r0	r1	r1	1/8
r1	r0	r1	1/8
r1	r1	r0	1/8
R0	R0	R0	1/8
R0	R1	R1	1/8
R1	R0	R1	1/8
R1	R1	R0	1/8

(a) $\Pr(x_1, x_2, y)$

$$\begin{aligned} \mathrm{I}(X_1, X_2 : Y) &= 2 \\ \mathrm{I}(X_1 : Y) &= 1 \\ \mathrm{I}(X_2 : Y) &= 1 \\ \mathrm{I}_{\min}\big(\{X_1, X_2\} : Y\big) &= 1 \\ \mathrm{I}_{\wedge}\big(\{X_1, X_2\} : Y\big) &= 1 \\ \mathrm{I}_{\alpha}\big(\{X_1, X_2\} : Y\big) &= 1 \end{aligned}$$

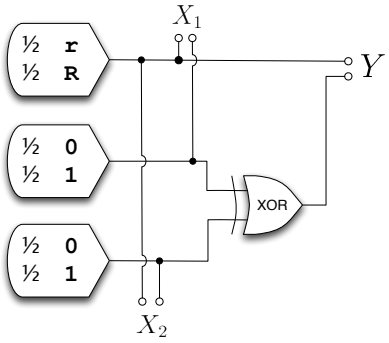

(b) circuit diagram

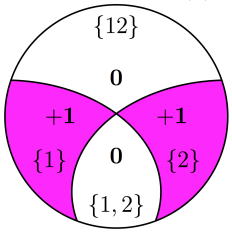

(c) Syn

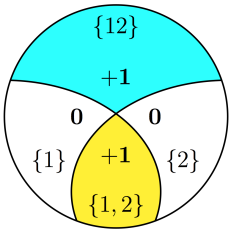

(d) $\Delta\mathrm{I}/\mathrm{I}_{\min}/\mathrm{I}_{\wedge}/\mathrm{I}_{\alpha}$

Fig. 3: Example RDNXOR. This is the canonical example of redundancy and synergy coexisting. $\mathrm{I}_{\min}$ and $\mathrm{I}_{\wedge}$ each reach the desired decomposition of one bit of redundancy and one bit of synergy. This example demonstrates $\mathrm{I}_{\wedge}$ correctly extracting the embedded redundant bit within X_1 and X_2.

IV. Previous candidate measures

In [13], the authors propose to use the following quantity, $\mathrm{I}_{\min}$, as the intersection information measure:

$$\begin{aligned} &\mathrm{I}_{\min}\left(X_1, \ldots, X_n : Y\right) \\ &\equiv \sum_{y \in Y} \Pr(y) \min_{i \in \{1,\ldots,n\}} \mathrm{I}(X_i : Y = y) \\ &= \sum_{y \in Y} \Pr(y) \min_{i \in \{1,\ldots,n\}} \mathrm{D}_{\mathrm{KL}}\Big[\Pr\big(X_i|y\big) \Big\| \Pr(X_i)\Big] \ , \end{aligned} \tag{3}$$

where D_{KL} is the Kullback-Leibler divergence.

Though $\mathrm{I}_{\min}$ is an intuitive and plausible choice for the intersection information, [9] showed that $\mathrm{I}_{\min}$ has counter-intuitive properties. In particular, $\mathrm{I}_{\min}$ calculates one bit of redundant information for example UNQ (Figure 2). It does this because each input shares one bit of information with the output. However, its quite clear that the shared informations are, in fact, different: X_1 provides the low bit, while X_2 provides the high bit. This led to the conclusion that $\mathrm{I}_{\min}$ *overestimates* the ideal intersection information measure by focusing only on *how much* information the inputs provide to the output. Another way to understand why $\mathrm{I}_{\min}$ overestimates redundancy in example UNQ is to imagine a hypothetical example where there are exactly two bits of unique information for every state $y \in Y$ and no synergy or redundancy. $\mathrm{I}_{\min}$ would calculate the redundancy as the minimum over both predictors which would be $\min[1, 1] = 1$ bit. Therefore $\mathrm{I}_{\min}$ would calculate 1 bit of redundancy even though by definition there was no redundancy but merely two bits of unique information.

X_1	X_2	Y	
0	0	0	1/4
0	1	0	1/4
1	0	0	1/4
1	1	1	1/4

(a) $\Pr(x_1, x_2, y)$

$$\begin{aligned} \mathrm{I}(X_1, X_2 : Y) &= 0.811 \\ \mathrm{I}(X_1 : Y) &= 0.311 \\ \mathrm{I}(X_2 : Y) &= 0.311 \\ \mathrm{I}_{\min}\big(\{X_1, X_2\} : Y\big) &= 0.311 \\ \mathrm{I}_{\wedge}\big(\{X_1, X_2\} : Y\big) &= 0 \\ \mathrm{I}_{\alpha}\big(\{X_1, X_2\} : Y\big) &= 0 \end{aligned}$$

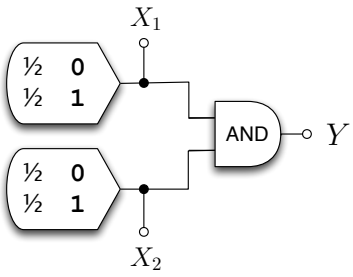

(b) circuit diagram

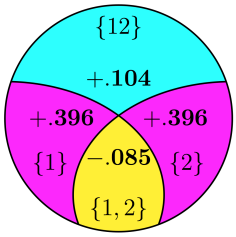

(c) $\Delta\mathrm{I}$

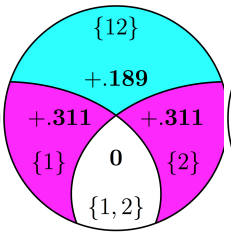

(d) $\mathrm{Syn}/\mathrm{I}_{\wedge}/\mathrm{I}_{\alpha}$

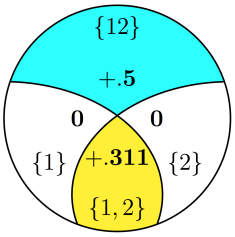

(e) $\mathrm{I}_{\min}$

Fig. 4: Example AND. $\Delta\mathrm{I}$ yields a redundant information of negative .085 bits. There are arguments for the redundant information being zero or positive, but thus far all that is agreed upon is that the redundant information is between $[0, 0.311]$ bits.

A candidate measure of synergy, $\Delta\mathrm{I}$, proposed in [14], leads to a negative value of redundant information for Example AND. Starting from $\Delta\mathrm{I}$ as a direct measure for synergistic information and then using eqs. (1) and (2) to derive the other terms, we get Figure 4c showing $\mathrm{I}_{\cap}\big(\{X_1, X_2\} : Y\big) \approx$

X_1 X_2	Y		
0 0	0	0.499	$\mathrm{I}(X_1, X_2 : Y) = 1$
0 1	0	0.001	$\mathrm{I}(X_1 : Y) = 1$
1 1	1	0.500	$\mathrm{I}(X_2 : Y) = 0.99$

(a) $\Pr(x_1, x_2, y)$

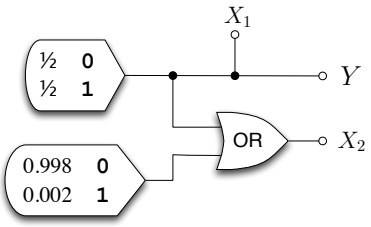

(b) circuit diagram

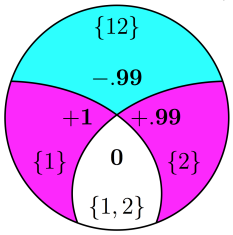

(c) $\mathrm{I}_\wedge$

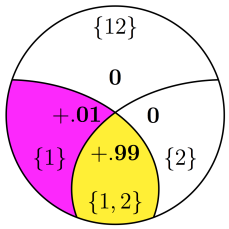

(d) Syn/Δ I/$\mathrm{I}_{\min}$/I_α

Fig. 5: Example ImperfectRdn. $\mathrm{I}_\wedge$ is blind to the noisy correlation between X_1 and X_2 and calculates zero redundant information. An ideal $\mathrm{I}_\cap$ measure would detect that all of the information X_2 specifies about Y is also specified by X_1 to calculate $\mathrm{I}_\cap\big(\{X_1, X_2\} : Y\big) = 0.99$ bits.

$-.085$ bits.

Another candidate measure of synergy, Syn [15], calculates zero synergy and redundancy for Example RdnXor, as opposed to the intuitive value of one bit of redundancy and one bit of synergy.

V. New candidate measures

A. The $\mathrm{I}_\wedge$ measure

Based on [16], we can consider a candidate intersection information as the maximum mutual information $\mathrm{I}(Q : Y)$ that some random variable Q conveys about Y, subject to Q being a function of each predictor $X_1, \ldots, X_n$. After some algebra, the leads to,

$$\begin{aligned} &\mathrm{I}_\wedge\big(\{X_1, \ldots, X_n\} : Y\big) \\ &\equiv \max_{\Pr(Q|Y)} \mathrm{I}(Q : Y) \\ &\quad \text{subject to } \forall i \in \{1, \ldots, n\} : \mathrm{H}\big(Q|X_i\big) = 0 \end{aligned} \quad , \tag{4}$$

which reduces to a simple expression in [12].

Example ImperfectRdn highlights the foremost shortcoming of $\mathrm{I}_\wedge$; $\mathrm{I}_\wedge$ does not detect "imperfect" or "lossy" correlations between X_1 and X_2. Instead, $\mathrm{I}_\wedge$ calculates zero redundant information, that $\mathrm{I}_\cap\big(\{X_1, X_2\} : Y\big) = 0$ bits. This arises from $\Pr(X_1 = 1, X_2 = 0) > 0$. If this were zero, ImperfectRdn reverts to being determined by the properties (**SR**) and the ($\mathbf{M_0}$) equality condition. Due to the nature of the common random variable, $\mathrm{I}_\wedge$ only sees the "deterministic" correlations between X_1 and X_2—add even an iota of noise between X_1 and X_2 and $\mathrm{I}_\wedge$ plummets to zero. This highlights a related issue with $\mathrm{I}_\wedge$; it is not continuous—an arbitrarily small change in the probability distribution can result in a discontinuous jump in the value of $\mathrm{I}_\wedge$.

Despite this, $\mathrm{I}_\wedge$ is useful stepping-stone, it captures what is inarguably redundant information (the common random variable). In addition, unlike earlier measures, $\mathrm{I}_\wedge$ satisfies (**TM**).

B. The I_α measure

Intuitively, we expect that if Q only specifies redundant information, that conditioning on any predictor X_i would vanquish all of the information Q conveys about Y. Noticing that $\mathrm{I}_\wedge$ *underestimates* the ideal $\mathrm{I}_\cap$ measure (i.e. it doesn't satisfy (**LP**)), we loosen the constraint $\mathrm{H}\big(Q|X_i\big) = 0$ in eq. (4), leading us to define the measure I_α:

$$\begin{aligned} &\mathrm{I}_\alpha\big(\{X_1, \ldots, X_n\} : Y\big) \\ &\equiv \max_{\Pr(Q|Y)} \mathrm{I}(Q : Y) \\ &\quad \text{subject to } \forall i \in \{1, \ldots, n\} : \mathrm{I}(Q, X_i : Y) = \mathrm{I}(X_i : Y) \ . \end{aligned} \tag{5}$$

$$\begin{aligned} &= \max_{\Pr(Q|Y)} \mathrm{I}(Q : Y) \\ &\quad \text{subject to } \forall i \in \{1, \ldots, n\} : \mathrm{H}(Q|X_i) = \mathrm{H}(Q|X_i, Y) \ . \end{aligned} \tag{6}$$

This measure obtains the desired values for the canonical examples in Section III. However, its implicit definition makes it more difficult to verify whether or not it satisfies the desired properties in Section III. Pleasingly, I_α also satisfies (**TM**).

We can also show (See Lemmas 1 and 2 in Appendix A) that

$$\begin{aligned} &0 \leq \mathrm{I}_\wedge\big(\{X_1, \ldots, X_n\} : Y\big) \leq \mathrm{I}_\alpha\big(\{X_1, \ldots, X_n\} : Y\big) \\ &\leq \mathrm{I}_{\min}\left(X_1, \ldots, X_n : Y\right) \ . \end{aligned} \tag{7}$$

While I_α satisfies previously defined canonical examples, we have found another example, shown in Figure 6, for which $\mathrm{I}_\wedge$ and I_α both calculate negative synergy. This example further complicates Example And by making the predictors mutually dependent.

$X_1\,X_2$	Y	
0 0	00	1/3
0 1	01	1/3
1 1	11	1/3

(a) $\Pr(x_1, x_2, y)$

$$\begin{aligned} \mathrm{I}(X_1, X_2 : Y) &= 1.585 \\ \mathrm{I}(X_1 : Y) &= 0.918 \\ \mathrm{I}(X_2 : Y) &= 0.918 \\ \mathrm{I}(X_1 : X_2) &= 0.252 \end{aligned}$$

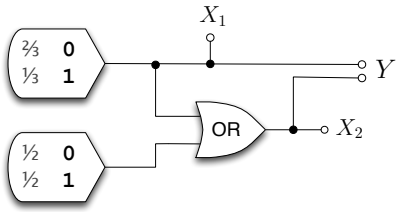

(b) circuit diagram

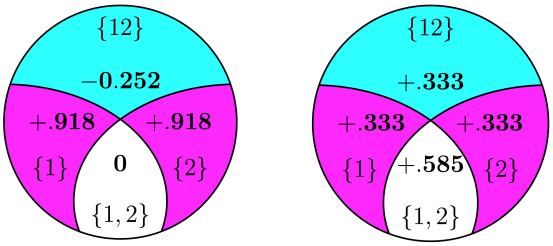

(c) $\mathrm{I}_\wedge/\mathrm{I}_\alpha$ (d) $\mathrm{I}_{\min}$

Fig. 6: Example Subtle.

VI. Conclusion

We have defined new measures for redundant information of predictor random variables regarding a target random variable. It is not clear whether it is possible for a single measure of synergy/redundancy to satisfy all previously proposed desired properties and canonical examples, and some of them are debatable. For example, a plausible measure of the "unique information" [17] and "union information" [9] yields an equivalent $\mathrm{I}_\cap$ measure that does not satisfy (**TM**). Determining whether some of these properties are contradictory is an interesting question for further work.

Acknowledgements. We thank Jim Beck, Yaser Abu-Mostafa, Edwin Chong, Chris Ellison, and Ryan James for helpful discussions.

References

[1] Schneidman E, Bialek W, II MB (2003) Synergy, redundancy, and independence in population codes. Journal of Neuroscience 23: 11539–53.

[2] Narayanan NS, Kimchi EY, Laubach M (2005) Redundancy and synergy of neuronal ensembles in motor cortex. The Journal of Neuroscience 25: 4207-4216.

[3] Balduzzi D, Tononi G (2008) Integrated information in discrete dynamical systems: motivation and theoretical framework. PLoS Computational Biology 4: e1000091.

[4] Anastassiou D (2007) Computational analysis of the synergy among multiple interacting genes. Molecular Systems Biology 3: 83.

[5] Lizier JT, Flecker B, Williams PL (2013) Towards a synergy-based approach to measuring information modification. CoRR abs/1303.3440.

[6] Gács P, Körner J (1973) Common information is far less than mutual information. Problems of Control and Informaton Theory 2: 149–162.

[7] Wyner AD (1975) The common information of two dependent random variables. IEEE Transactions in Information Theory 21: 163–179.

[8] Kumar GR, Li CT, Gamal AE (2014) Exact common information. CoRR abs/1402.0062.

[9] Griffith V, Koch C (2012) Quantifying synergistic mutual information. CoRR abs/1205.4265.

[10] Harder M, Salge C, Polani D (2013) Bivariate measure of redundant information. Phys Rev E 87: 012130.

[11] Bertschinger N, Rauh J, Olbrich E, Jost J (2012) Shared information – new insights and problems in decomposing information in complex systems. CoRR abs/1210.5902.

[12] Griffith V, Chong EKP, James RG, Ellison CJ, Crutchfield JP (2013) Intersection information based on common randomness. CoRR abs/1310.1538.

[13] Williams PL, Beer RD (2010) Nonnegative decomposition of multivariate information. CoRR abs/1004.2515.

[14] Nirenberg S, Latham PE (2003) Decoding neuronal spike trains: How important are correlations? Proceedings of the National Academy of Sciences 100: 7348–7353.

[15] Schneidman E, Still S, Berry MJ, Bialek W (2003) Network information and connected correlations. Phys Rev Lett 91: 238701-238705.

[16] Wolf S, Wullschleger J (2004) Zero-error information and applications in cryptography. Proc IEEE Information Theory Workshop 04: 1–6.

[17] Bertschinger N, Rauh J, Olbrich E, Jost J, Ay N (2013) Quantifying unique information. CoRR abs/1311.2852.

Appendix

Proof I_α does not satisfy (LP). Proof by counter-example Subtle (Figure 6). For $\mathrm{I}\big(Q : Y | X_1\big) = 0$, then Q must not distinguish between states of $Y = \mathtt{00}$ and $Y = \mathtt{01}$ (because X_1 does not distinguish between these two states). This entails that $\Pr\big(Q | Y = \mathtt{00}\big) = \Pr\big(Q | Y = \mathtt{01}\big)$. By symmetry, likewise for $\mathrm{I}\big(Q : Y | X_2\big) = 0$, Q must be distinguish between states $Y = \mathtt{01}$ and $Y = \mathtt{11}$. Altogether, this entails that $\Pr\big(Q | Y = \mathtt{00}\big) = \Pr\big(Q | Y = \mathtt{01}\big) = \Pr\big(Q | Y = \mathtt{11}\big)$, which then entails, $\Pr\big(q | y_i\big) = \Pr\big(q | y_j\big) \quad \forall q \in Q, y_i \in Y, y_j \in Y$, which is only achievable when $\Pr(q) = \Pr\big(q | y\big) \quad \forall q \in Q, y \in Y$. This makes $\mathrm{I}(Q : Y) = 0$, therefore for example Subtle, $\mathrm{I}_\alpha\big(\{X_1, X_2\} : Y\big) = 0$.

Lemma 1. *We have* $\mathrm{I}_\wedge\big(\{X_1, \ldots, X_n\} : Y\big) \leq \mathrm{I}_\alpha\big(\{X_1, \ldots, X_n\} : Y\big)$.

Proof. We define a random variable $Q' = X_1 \wedge \cdots \wedge X_n$. We then plugin Q' for Q in the definition of I_α. This newly plugged-in Q satisfies the constraint $\forall i \in \{1, \ldots, n\}$ that

$\mathrm{I}\big(Q\!:\!Y|X_i\big) = 0$. Therefore, Q' is always a possible choice for Q, and the maximization of $\mathrm{I}(Q\!:\!Y)$ in I_α must be at least as large as $\mathrm{I}\big(Q'\!:\!Y\big) = \mathrm{I}_\wedge\big(\{X_1,\ldots,X_n\}\!:\!Y\big)$. □

Lemma 2. *We have* $\mathrm{I}_\alpha\big(\{X_1,\ldots,X_n\}:Y\big) \leq \mathrm{I}_{\min}\left(X_1,\ldots,X_n:Y\right)$

Proof. For a given state $y \in Y$ and two arbitrary random variables Q and X, given $\mathrm{I}\big(Q\!:\!y|X\big) = \mathrm{D}_{\mathrm{KL}}\Big[\Pr\big(QX|y\big)\Big\|\Pr\big(Q|X\big)\Pr\big(X|y\big)\Big] = 0$, we show that, $\mathrm{I}(Q\!:\!y) \leq \mathrm{I}(X\!:\!y)$,

$$\mathrm{I}(X\!:\!y) - \mathrm{I}(Q\!:\!y) = \sum_{x\in X}\Pr\big(x|y\big)\log\frac{\Pr\big(x|y\big)}{\Pr(x)} - \sum_{q\in Q}\Pr\big(q|y\big)\log\frac{\Pr\big(q|y\big)}{\Pr(q)} \geq 0\,.$$

Generalizing to n predictors $X_1,\ldots,X_n$, the above shows that that the maximum $\mathrm{I}(Q\!:\!y)$ under constraint $\mathrm{I}\big(Q\!:\!y|X_i\big)$ will always be less than $\min_{i\in\{1,\ldots,n\}}\mathrm{I}(X_i\!:\!y)$, which completes the proof. □

Lemma 3. *Measure* $\mathrm{I}_{\min}$ *satisfies desired property Strong Monotonicity,* $(\mathbf{M_1})$.

Proof. Given $\mathrm{H}\big(Y|Z\big) = 0$, then the specific-surprise $\mathrm{I}(Z\!:\!y)$ yields,

$$\begin{aligned}\mathrm{I}(Z\!:\!y) &\equiv \mathrm{D}_{\mathrm{KL}}\Big[\Pr\big(Z|y\big)\Big\|\Pr(Z)\Big]\\ &= \sum_{z\in Z}\Pr\big(z|y\big)\log\frac{\Pr\big(z|y\big)}{\Pr(z)}\\ &= \sum_{z\in Z}\Pr\big(z|y\big)\log\frac{1}{\Pr(y)}\\ &= \log\frac{1}{\Pr(y)}\,.\end{aligned}$$

Given that for an arbitrary random variable X_i, $\mathrm{I}(X_i\!:\!y) \leq \log\frac{1}{\Pr(y)}$. As $\mathrm{I}_{\min}$ takes only uses the $\min_i \mathrm{I}(X_i\!:\!y)$, the minimum is invariant under adding any predictor Z such that $\mathrm{H}\big(Y|Z\big) = 0$. Therefore, measure $\mathrm{I}_{\min}$ satisfies property $(\mathbf{M_1})$. □